

Europäisches Patentamt  
European Patent Office  
Office européen des brevets



EP 1 014 264 A2

(12)

## EUROPEAN PATENT APPLICATION

(43) Date of publication:  
28.06.2000 Bulletin 2000/26

(51) Int CL: G06F 9/46

(21) Application number: 99309854.0

(22) Date of filing: 07.12.1999

(34) Designated Contracting States:  
AT BE CH CY DE DK ES FI FR GB GR IE IT LU LU  
MC NL PT SE  
Designated Extension States:  
AL LT LV MK RO SI

(72) Inventor: Schmidt, Rene W.  
Cupertino, California 95014 (US)

(30) Priority: 07.12.1998 US 206810

(74) Representative: Browne, Robin Forsythe, Dr.  
Urquhart-Dykes & Lord  
Tower House  
Merton Way  
Leeds LS2 8PA West Yorkshire (GB)

(71) Applicant: SUN MICROSYSTEMS, INC.  
Palo Alto, California 94043 (US)

(54) A system and method for facilitating safepoint synchronization in a multithreaded computer system

(57) The present invention is a system and method for facilitating a global safepoint operation in a multi-threaded computer system. According to an embodiment of the present invention, each thread keeps track of its safepoint regions by maintaining a variable (a flag) which indicates a state, such as whether the current region of the thread is safe, unsafe, or transitional. In this manner, it can be determined whether a thread is cur-

rently in a safepoint region without suspending the thread. When a thread is currently in a safepoint region, the thread can continue to operate while a global safepoint operation, such as garbage collection is being performed. When the thread begins to transition out of the safe region, it moves into a transitional region. The transitional region automatically blocks the transition into the non-safe region to assure that the safepoint operation occurs in a safe region.

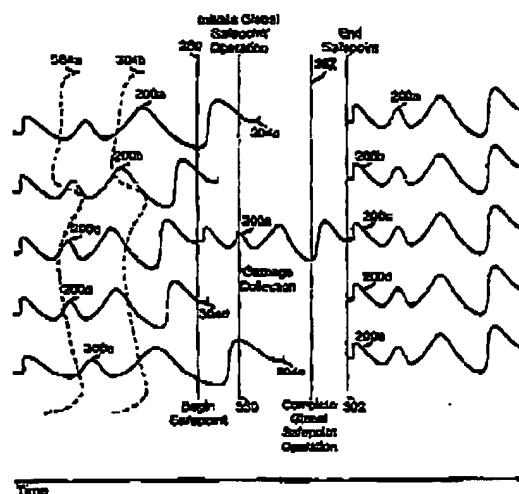


FIG. 3

EP 1 014 264 A2

**Description****FIELD OF THE INVENTION**

[0001] The present invention relates to the facilitation of a global safepoint operation in a multithreaded computer system environment. In particular, the present invention relates to facilitating a global safepoint operation while avoiding suspending all threads during safepoint synchronization.

**BACKGROUND OF THE INVENTION**

[0002] In computer science, the term "multitasking" refers to an operating system's ability to support multiple processes simultaneously. A process is typically a program in execution. Support for multiple processes is necessary in applications where several computations proceed in parallel. In time-sharing systems, multiple users share a single computer system and all processes created by them should, at least in principle, execute simultaneously. Real time systems that control multiple devices also need to support multiple processes. For instance, an avionics computer on board an airplane runs processes for monitoring the engines, updating the flight instruments, processing radar signals, and keeping the airplane on course. Batch operating systems depend on multitasking for overlapping computation with input/output (I/O) operations. For instance, when a process performs I/O, the operating system may run another process to avoid idling the central processor for long periods of time.

[0003] One type of process is known as primitive process or threads. A single process may also include multiple threads. Typically, the simplest way to execute multiple threads simultaneously is to assign each thread to its own processor in a multiprocessor system. If the number of threads exceeds the number of processors, then processors must typically be multiplexed among threads. By switching a processor rapidly from one thread to the next, it appears to the observer as if all threads are making progress, even if the processor can execute only one instruction at a time. While processor multiplexing typically implements only quasi-parallelism, peripheral devices can provide true parallelism even if the computer system contains only a single, central processor. Peripheral devices can be regarded as specialized processors that operate concurrently with the central processor. A device runs a single process specialized, for example, for printing a line or writing a disk block. The device receives commands from a device driver process that itself runs on the central processor. After a device driver has issued a command to a device, the driver waits for a completion signal. During this wait, the main processor typically switches its attention from the device driver to other threads.

[0004] While maintaining multiple threads, the computer system may need to perform global operations

which require synchronization or control of all or a group of threads at a given time. An example of such a global operation is garbage collection.

[0005] Many programming languages and systems provide for dynamic as well as static allocation of memory storage to abstract data objects. The performance of these systems relies on their ability to reclaim and reuse storage for dynamically allocated objects after they are no longer needed by the executed program.

[0006] Some language systems require programmers to return unneeded objects (garbage) to the memory system explicitly. Although this permits precise and efficient recycling of storage when performed carefully, it often results in objects being recycled prematurely or being forgotten and thus lost to the system. Other systems reclaim abandoned objects automatically through a process called "garbage collection". Reclaiming storage automatically in this way is both a convenience to the programmer and a means for insuring every object's storage is recycled correctly.

[0007] Garbage collection typically occurs in two phases: identifying unneeded objects, and then making their storage available for reallocation. An object in a program is needed, or live at a given time if the program might access the object in the future; otherwise it is dead. In practice, garbage collectors typically consider an object to be dead only if the program has abandoned all pointers to it, making future access impossible.

[0008] For a global operation, such as garbage collection, all the locations of objects and all locations of reference pointers typically need to be known to perform such a global operation. In a multithreaded environment, a "stop" instruction is typically sent to all threads prior to performing a global operation, such as garbage collection. Then it needs to be determined whether each thread is in a 'safe' region or an 'unsafe' region. A safe region is a region of code through which a thread is processing, wherein pointers are not being manipulated. An unsafe region is a region of the code through which the thread is processing wherein pointers may be manipulated. Conventionally, all threads are typically suspended ("thread suspends") in order to evaluate each thread and determine if each thread is in a safe or unsafe region. If a thread is in an unsafe region, then the thread operation is resumed and is stopped later to attempt to suspend it at a safe region.

[0009] Although thread suspends can be effective in some systems, typically, the more concurrent threads that are in use at any given time, the slower the suspend process progresses. Accordingly, suspend processes can be very expensive to perform on many of today's advanced processors. Additionally, many threads are in a situation such that they do not need to be suspended. For example, when a thread is in a safe region, then that particular thread causes no problems for the performance of a global safepoint operation. However, conventionally, all the threads are suspended in order to evaluate and determine whether each thread is in a safe or

[0005] It would be desirable to facilitate a global safe-point operation, without the need to suspend all threads. The present invention addresses such a need.

## **SUMMARY OF THE INVENTION**

[0010] The present invention is a system and method for facilitating a global safepoint operation in a multi-threaded computer system. According to an embodiment of the present invention, each thread keeps track of its safepoint regions by maintaining a variable which indicates a state such as whether the current region of the thread is safe, unsafe or transitional. In this manner, it can be determined whether a thread is currently in a safepoint region without suspending the thread. When a thread is currently in a safepoint region, the thread can continue to operate while a global safepoint operation, such as garbage-collection is being performed. When the thread begins to transition out of the safe region, it moves into a transitional region. The transitional region automatically blocks the transition into the non-safe region to assure that the safepoint operation occurs in a safe region.

[0011] According to an embodiment of the present invention, one thread may begin a safepoint procedure after it acquires a safepoint lock. The safepoint lock blocks other threads from attempting to perform the safepoint procedure at the same time. Prior to initiating a safepoint operation, the thread attempting to initiate the safepoint operation will wait until all other threads reach a safe region in each of their own code.

[0012] A method according to an embodiment of the present invention for facilitating a safepoint operation in a multi-threaded computer system is presented. A first thread acquires a safepoint; and a determination is made as to whether at least one of the remaining threads is in a safe region, requiring the suspension of all of the threads. A safepoint operation is then initiated.

[0013] In another aspect of the invention, a method according to an embodiment of the present invention for facilitating a safepoint operation in a computer system including at least one thread is also presented. The method comprising determining whether a thread is in a safe region, the determination being made without suspending the thread; continuing thread operation if the thread is in the safe region; and initiating a safepoint operation while the thread is continuing to operate.

[0014] In yet another aspect of the invention, a method according to an embodiment of the present invention for facilitating a safepoint operation in a computer system including at least one thread is also presented. The method comprising providing a state associated with a thread; determining if the state indicates a transition; and blocking an operation of the thread if the state indicates a predetermined type of transition and if a safepoint operation has been initiated.

[0015] In yet another aspect of the invention, a system

according to an embodiment of the present invention for facilitating a safepoint operation in a computer system including at least one thread is also presented. The system comprising a variable indicating whether a thread in operation is in a safe region; and a processor configured to read the flag, the processor also being configured to determine if the thread is in a safe region without suspending the thread; the processor also being configured to initiate a safepoint operation if the thread is in a safe region.

**BRIEF DESCRIPTION OF THE DRAWINGS**

[0016] Figure 1 is a block diagram of a computer system suitable for implementing the present invention.

[0017] Figure 2 is an illustration of a thread transitioning through safe and unsafe regions of a computer code. [0018] Figure 3 is an illustration of threads progressing through a global snapshot operation according to an

[0019] Figure 4 is a flow diagram of a transition within a thread according to an embodiment of the present invention.

[0020] Figure 5 is a flow diagram of a method for 25 blocking a thread according to an embodiment of the present invention.

[0021] Figure 6 is a flow diagram of a method for implementing a notify block according to an embodiment of the present invention.

30 [0022] Figure 7A - 7E are flow diagrams of a method for setpoint synchronization according to an embodiment of the present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0023] The following description is presented to enable one of ordinary skill in the art to make and to use the invention and is provided in the context of a patent application and its requirements. Various modifications to the preferred embodiments will be readily apparent to those skilled in the art and the generic principles herein may be applied to other embodiments. Thus, the present invention is not intended to be limited to the embodiment shown but is to be accorded the widest scope consistent with the principles and features described herein.

[0024] Figure 1 is a block diagram of a general purpose computer system 100 suitable for carrying out the processing in accordance with one embodiment of the present invention. Figure 1 illustrates one embodiment of a general purpose computer system. Other computer system architectures and configurations can be used for carrying out the processing of the present invention.

50 Computer system 100, made up of various subsystems described below, includes at least one microprocessor subsystem (also referred to as a central processing unit, or CPU) 102. That is, CPU 102 can be implemented by

55

a single-chip processor or by multiple processors. CPU 102 is a general purpose digital processor which controls the operation of the computer system 100. Using instructions retrieved from memory 110, the CPU 102 controls the reception and manipulation of input data, and the output and display of data on output devices. [0025] CPU 102 is coupled bi-directionally with memory 110 which can include a first primary storage, typically a random access memory (RAM), and a second primary storage area, typically a read-only memory (ROM). As is well known in the art, primary storage can be used as a general storage area and as scratch-pad memory, and can also be used to store input data and processed data. It can also store programming instructions and data, in the form of data objects and text objects, in addition to other data and instructions for processes operating on CPU 102. Also as well known in the art, primary storage typically includes basic operating instructions, program code, data and objects used by the CPU 102 to perform its functions. Primary storage devices 110 may include any suitable computer-readable storage media, described below, depending on whether, for example, data access needs to be bi-directional or uni-directional. CPU 102 can also directly and very rapidly retrieve and store frequently needed data in a cache memory (not shown).

[0026] A removable mass storage device 112 provides additional data storage capacity for the computer system 100, and is coupled either bi-directionally or uni-directionally to CPU 102. For example, a specific removable mass storage device commonly known as a CD-ROM typically passes data uni-directionally to the CPU 102, whereas a floppy disk can pass data bi-directionally to the CPU 102. Storage 112 may also include computer-readable media such as magnetic tape, flash memory, signals embodied on a carrier wave, PC-CARDS, portable mass storage devices, holographic storage devices, and other storage devices. A fixed mass storage 120 can also provide additional data storage capacity. The most common example of mass storage 120 is a hard disk drive. Mass storage 112, 120 generally store additional programming instructions, data, and the like that typically are not in active use by the CPU 102. It will be appreciated that the information retained within mass storage 112, 120 may be incorporated, if needed, in standard fashion as part of primary storage 110 (e.g. RAM) as virtual memory.

[0027] In addition to providing CPU 102 access to storage subsystems, bus 114 can be used to provide access other subsystems and devices as well. In the described embodiment, these can include a display monitor 118, a network interface 116, a keyboard 104, and a pointing device 106, as well as an auxiliary input/output device interface, a sound card, speakers, and other subsystems as needed. The pointing device 106 may be a mouse, stylus, track ball, or tablet, and is useful for interacting with a graphical user interface.

[0028] The network interface 116 allows CPU 102 to

be coupled to another computer, computer network, or telecommunications network using a network connection as shown. Through the network interface 116, it is contemplated that the CPU 102 might receive information, e.g., data objects or program instructions, from another network, or might output information to another network in the course of performing the above-described method steps. Information, often represented as a sequence of instructions to be executed on a CPU, may be received from and outputted to another network, for example, in the form of a computer data signal embodied in a carrier wave. An interface card or similar device and appropriate software implemented by CPU 102 can be used to connect the computer system 100 to an external network and transfer data according to standard protocols. That is, method embodiments of the present invention may execute solely upon CPU 102, or may be performed across a network such as the Internet, intranet networks, or local area networks, in conjunction with a remote CPU that shares a portion of the processing. Additional mass storage devices (not shown) may also be connected to CPU 102 through network interface 116.

[0029] An auxiliary I/O device interface (not shown) can be used in conjunction with computer system 100. The auxiliary I/O device interface can include general and customized interfaces that allow the CPU 102 to send and, more typically, receive data from other devices such as microphones, touch-sensitive displays, transducer card readers, tape readers, voice or handwriting recognizers, biometrics readers, cameras, portable mass storage devices, and other computers.

[0030] In addition, embodiments of the present invention further relate to computer storage products with a computer readable medium that contain program code for performing various computer-implemented operations. The computer-readable medium is any data storage device that can store data which can thereafter be read by a computer system. The media and program code may be those specially designed and constructed for the purposes of the present invention, or they may be of the kind well known to those of ordinary skill in the computer software arts. Examples of computer-readable media include, but are not limited to, all the media mentioned above; magnetic media such as hard disks, floppy disks, and magnetic tape; optical media such as CD-ROM disks; magneto-optical media such as floptical disks; and specially configured hardware devices such as application-specific integrated circuits (ASICs), programmable logic devices (PLDs), and ROM and RAM devices. The computer-readable medium can also be distributed as a data signal embodied in a carrier wave over a network of coupled computer systems so that the computer-readable code is stored and executed in a distributed fashion. Examples of program code include both machine code, as produced, for example, by a compiler, or files containing higher level code that may be executed using an interpreter.

[0021] The computer system shown in Fig. 1 is but an example of a computer system suitable for use with the invention. Other computer systems suitable for use with the invention may include additional or fewer subsystems. In addition, bus 114 is illustrative of any interconnection scheme serving to link the subsystems. Other computer architectures having different configurations of subsystems may also be utilized.

[0032] The present invention is a system and method for facilitating a global safepoint operation in a multi-threaded computer system. According to an embodiment of the present invention, each thread keeps track of its safepoint regions by maintaining a variable which indicates a state, such as whether the current region of the thread is safe, unsafe, or transitional. In this manner, it can be determined whether a thread is currently in a safepoint region without suspending the thread. When a thread is currently in a safepoint region, the thread can continue to operate while a global safepoint operation, such as garbage-collection is being performed. When the thread begins to transition out of the safe region, it moves into a transitional region. The transitional region automatically blocks the transition into the non-safe region to assure that the safepoint operation occurs in a safe region.

[0033] According to an embodiment of the present invention, one thread may begin a safepoint procedure after it acquires a safepoint lock. The safepoint lock blocks other threads from attempting to perform the safepoint procedure at the same time. Prior to initiating a safepoint operation, the thread attempting to initiate the safepoint operation will wait until all other threads reach a safe region in each of their own code.

[0034] Figure 2 is an illustration of a thread transitioning between safe and unsafe regions of a computer code. Figure 2 shows thread 200 transitioning from a safe region 202A to an unsafe region 204 via a transition region 206A. The transition region 206A is a transition from a safe region to an unsafe region. The thread 200 is also shown to transition from the unsafe region 204 to a safe region 202B via a transition region 206B, wherein the transition region 206B is a transition from an unsafe region to a safe region.

[0035] If thread 200 is in a safe region, such as safe regions 202A and 202B, a global safepoint operation, such as garbage collection, may proceed while thread 200 continues to operate. When thread 200 reaches the transition region 206A, transitioning from a safe region to an unsafe region, thread 200 may be blocked from continuing into the unsafe region 204 to insure that thread 200 stays in a safe region during the global safepoint operation.

[0035] According to an embodiment of the present invention, if thread 200 is in unsafe region 204, then thread 200 is allowed to continue through unsafe region 204 and the global safepoint operation is not initiated. As previously mentioned, unsafe regions are regions of code through which a thread is passing wherein pointers

may be manipulated. Examples of unsafe regions include virtual machine code, Java™ code, and code which is unidentified. The Java™ code may include compiler code and interpretive code.

6 [0037] When thread 200 enters the transition region 205B, transitioning from an unsafe area to a safe area, a waiting counter is decremented. Further details of the waiting counter will later be discussed in conjunction with the remaining figures, such as Figure 8.

[0038] Figure 3 is an illustration of a set of threads progressing through a global assignment operating ac-

progressing through a global computer operation according to an embodiment of the present invention. Figure 8 shows a set of threads 200A-200E, each progressing through their own set of computer codes. Points 304A and 304B represent state transitions for each thread 200A-200E. For example, point 304A for thread 200A may be a transition from a safe region to an unsafe region, while point 304A for thread 200B may be a transition from an unsafe region to a safe region. Likewise, point 304B for thread 200D may be a transition from an unsafe region to a safe region, while point 204B for thread 200E may be a transition from a safe region to an unsafe region.

[p029] Assume that it is desirable to perform a global safepoint operation, such as garbage collection, for thread 200C. The global safepoint operation is expected to be programmed such that it is code which is considered to be safe for use. Assume that the global safepoint operation is to be performed by thread 200C.

90 ered to be a safe region. Accordingly, when thread 200C initiates begin safepoint 200, it is assumed that thread 200C is moving through a safe region of code. Thus, thread 200C's state flag should not need to be checked.

[0040] Thread 200C then acquires a safepoint lock at

begin safepoint 300. Since thread 200C is the thread which initiates the safepoint operation, it is often herein referred to as the initiating thread. The safepoint lock

ensures that thread 200C is the only thread allowed to perform the global safepoint operation, thus avoiding simultaneous attempts at performing global safepoint op-

40 erations by different threads. Additionally, since the remaining threads 200A-200B, 200D-200E cannot acquire the safepoint lock, an attempt to acquire the safepoint lock by the remaining threads 200A-200B, 200D-200E will fail.

45 200E will block further progress by these threads. Accordingly, when the remaining threads 200A-200B, 200D-200E move into a transition region from a safe region, they may be blocked from progressing further into an unsafe region by causing these threads to attempt to acquire the safepoint lock.

[0041] After acquiring the endpoint lock, thread 200C then waits until every other thread 200A, 200B, 200D, 200E, has reached a safe region within their code. When a remaining thread 200A-200B, 200D-200E transitions from an unsafe region (such as region 204 of Figure 2) to a safe region (such as region 202B of Figure 2), a waiting counter is decremented to indicate that this thread is one less thread for thread 200C to wait upon. If any of the remaining threads 200A, 200B, 200D, 200E are in safe regions, they continue to progress until each

of those threads runs into a transition region from a safe region to an unsafe region. When a thread moves into such a transition, it is blocked to keep it from progressing out of the safe region and into an unsafe region. When all the threads 200A-200E are in safe regions, or are blocked at transitions from a safe region to an unsafe region, then the global safepoint operation, such as garbage collection, may be initiated.

[0042] For example, assume that at point 304B, thread 200A transitions into a safe region. Then thread 200A may continue to operate after begin safepoint 300 until it reaches a transition state 304C. When begin safepoint 300 is initiated, thread 200A's state flag will be checked to determine the type of region thread 200A is moving through. In this example, thread 200A's state flag will indicate that it is in a safe region. At transition state 304C, thread 200A will be blocked to ensure that it does not move into an unsafe region during the global safepoint operation.

[0043] Assume also that thread 200B transitions into an unsafe region at point 304B. When begin safepoint 300 is initiated, thread 200B's state flag is checked to determine what type of region thread 200B is moving through. In this example, thread 200B's state flag will indicate that it is in an unsafe region. Thread 200B is then suspended and handled such that a safepoint operation may be performed. Further details of unsafe region handling will later be discussed in conjunction with Figures 7B-7D.

[0044] If thread 200D is in a transition from a safe region to an unsafe region at the time thread 200D's state flag is checked after begin safepoint 300 is initiated. Then the progress of thread 200D is blocked to keep thread 200D from moving into an unsafe region. If thread 200E is in a transition from an unsafe region to a safe region at the time thread 200D's state flag is checked (after begin safepoint 300 is initiated), then thread 200E is allowed to continue its progress. If thread 200E reaches another transition state indicating a transition into an unsafe region, then it is blocked at point 304E.

[0045] When all the threads 200A-200E are in safe regions, then the safepoint operation may be initiated at point 350. When the global safepoint operation is complete at point 352, then a variable is set such that threads 200A-200E are no longer blocked at transitions, the safepoint lock is released by the initiating thread 200C, and any threads which have been suspended are then unsuspended. When these tasks are complete, then the safepoint procedure is ended at point 302.

[0046] Figures 4-6 describe events which occur for each thread at checkpoints 304A - 304D of Figure 3, while Figures 7A-7E describe events which occur at begin and end safepoints 300 and 302 of Figure 3.

[0047] Figure 4 is a flow diagram of a transition between program code regions within a thread according to an embodiment of the present invention. A thread local state variable is set to an incoming transitory state in step 400. The thread local state variable may be a

flag indicating the state of the region through which the thread is passing. For example, if a thread was transitioning from a safe region to an unsafe region, the transition region may be indicated by a number which sets the thread local state variable to indicate that it is a transition from a safe region to an unsafe region. Examples of other states which may be indicated by the thread local state variable include safe region, unsafe region, transition from a safe region to an unsafe region, and a transition from an unsafe region to a safe region.

[0048] It is then determined whether a variable, called "block at transition" variable, is set to true in step 402. The "block at transition" variable indicates whether a thread should be blocked at a transition region.

[0049] Note that a thread block is simpler and easier to perform than a thread suspend. When a thread is blocked, it simply does not proceed on to the next state or region, whereas when threads are suspended, the operating system typically must perform the suspend operation. Accordingly, it is simpler and less expensive to block a thread than to suspend the thread.

[0050] If "block at transition" variable is set to true in step 402, then a thread blocker is called in step 404. Details of the thread blocker will be discussed later in conjunction with figure 5. Thereafter, the thread local state variable is set to the next state which is entered by the thread in step 406.

[0051] If "block at transition" variable is not set to true in step 402, then thread local state variable is set to the next state (region) entered by the thread in step 408, and the transition is completed in step 408.

[0052] Figure 5 is a flow diagram of a method for blocking a thread according to an embodiment of the present invention. The flow diagram of Figure 5 occurs when the thread blocker is called in step 404 of Figure 4.

[0053] A notify block is implemented in step 500. Details of the implementation of the notify block will later be discussed in conjunction with Figure 5.

[0054] The thread local state variable is then set to a "blocked" state in step 502. A safepoint lock is then acquired in step 504. If the safepoint lock may not be immediately acquired (because the initiating thread has already acquired it), then it waits until the safepoint lock may be acquired. Waiting to acquire the safepoint lock keeps the thread in a "blocked" state by holding its position until the safepoint lock is acquired. The "blocked" state indicates to the initiating thread that this thread is in a safe region. The safepoint lock may only be acquired after the thread holding the safepoint lock releases it. Since the thread holding the safepoint lock will not release the safepoint lock until the global safepoint operation is completed, this thread is blocked throughout the global safepoint operation by waiting to acquire the safepoint lock.

[0055] The acquisition of a safepoint lock allows the initiating thread, such as thread 200C of Figure 3, to progress beyond the begin safepoint 300 and have the desired safepoint operation executed upon it, without

15

EP 1 014 264 A2

16

5. A method for facilitating a safepoint operation in a multi-threaded computer system, the method comprising:

determining whether a selected thread is in a safe region, the determination being made without suspending the selected thread; continuing thread operation if the selected thread is in the safe region; and initiating a safepoint operation while the selected thread is continuing to operate.

6. The method of claim 5, further detecting when the selected thread moves into a transition region from the safe region.

7. The method of claim 6, wherein the selected thread is blocked when the selected thread is in the transition region.

8. The method of any of claims 5-7, wherein the selected thread is blocked from moving out of the safe region.

9. The method of claim 8, wherein the selected thread is blocked by requiring that the selected thread acquire a safepoint lock before moving out of the safe region.

10. The method of any of claims 5-9, further comprising suspending the selected thread if the thread is in an unsafe region.

11. The method of claims 5-10, further comprising continuing the selected thread operation if the selected thread is in a transition from an unsafe region to a safe region, and initiating the safepoint operation only when the selected thread is in a safe region.

12. The method of any of claims 5-11, further comprising blocking the thread operation if the selected thread is in a transition from a safe region to an unsafe region.

13. A method for facilitating a safepoint operation in a multi-threaded computer system, the method comprising:

providing a state variable associated with a selected thread; determining if the state variable indicates that the selected thread is in a transition; and blocking an operation of the selected thread if the state indicates a predetermined type of transition and if a safepoint operation has been initiated.

14. The method of claim 13, wherein the predetermined type of transition is a transition from a safe region to an unsafe region.

15. The method of claim 13 or 14, further comprising decrementing a waiting counter if the state indicates a transition, wherein the transition is a transition from an unsafe region to a safe region.

16. The method of any of claims 13-15, wherein the state variable is a flag.

17. A system for facilitating a safepoint operation in a computer system including at least one thread, the system comprising:

means for acquiring a safepoint lock by a first thread, the first thread being one of a plurality of threads; means for determining if at least one of the plurality of threads is in a safe region, wherein the determination is made without suspending the at least one of the plurality of threads; and means for initiating a safepoint operation.

18. A system for facilitating a safepoint operation in a multi-threaded computer, the system comprising:

a data structure including a plurality of state variables, each state variable being arranged to whether an associated thread is in a safe region of operation.

19. The system of claim 18, further comprising a processor configured to read the plurality of state variables, the processor also being configured to determine if a thread is in a safe region without suspending the thread; the processor also being configured to initiate a safepoint operation if the thread is in a safe region.

20. A computer program product for facilitating a safepoint operation in a computer system including at least one thread, the computer program product comprising:

computer code acquiring a safepoint lock by a first thread, the first thread being one of a plurality of threads; computer code determining if at least one of the plurality of threads is in a safe region, wherein the determination is made without suspending the at least one of the plurality of threads; computer code initiating a safepoint operation; and a computer readable medium that stores the computer codes.

21. The computer program product of claim 20, wherein the computer readable medium is selected from the

17

EP 1 014 264 A2

18

group consisting of CD-ROM, floppy disk, tape,  
flash memory, system memory, hard drive, and data  
signal embodied in a carrier wave.

5

10

15

20

25

30

35

40

45

50

55

10

EP 1 014 264 A2

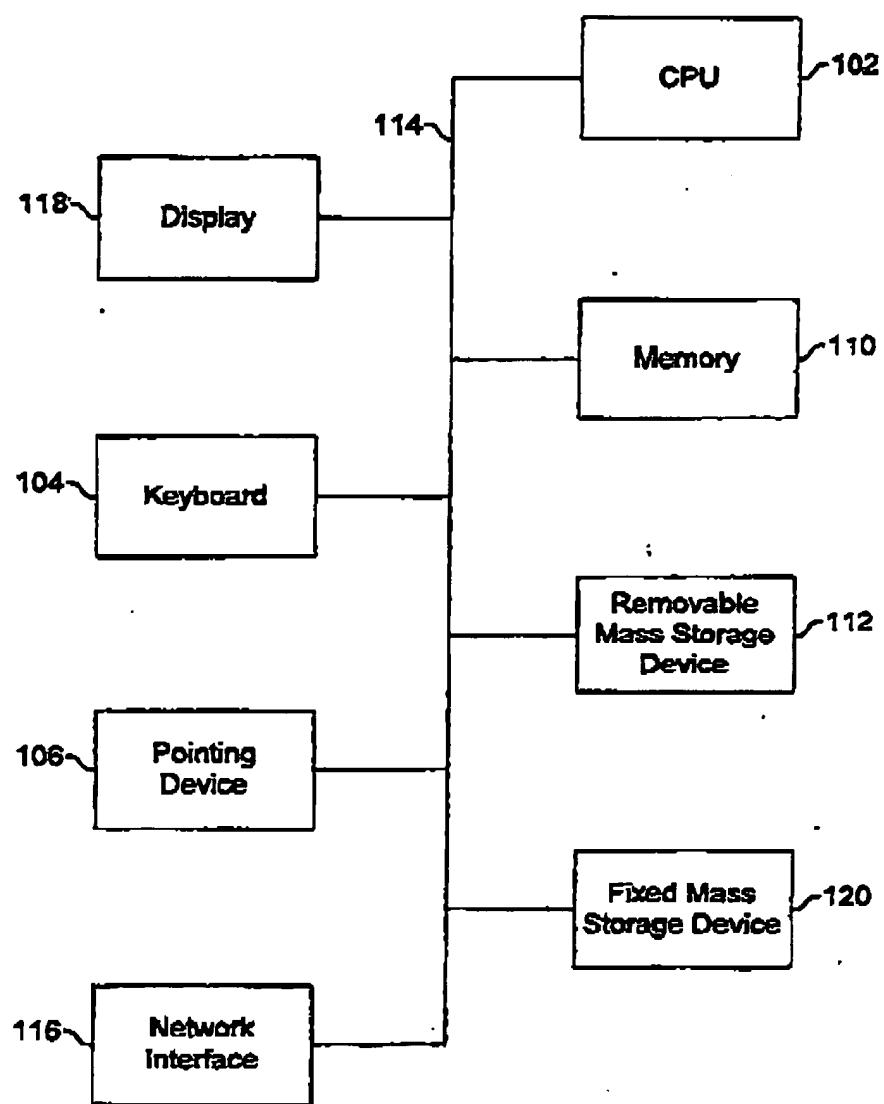


FIG. 1

EP 1 014 264 A2

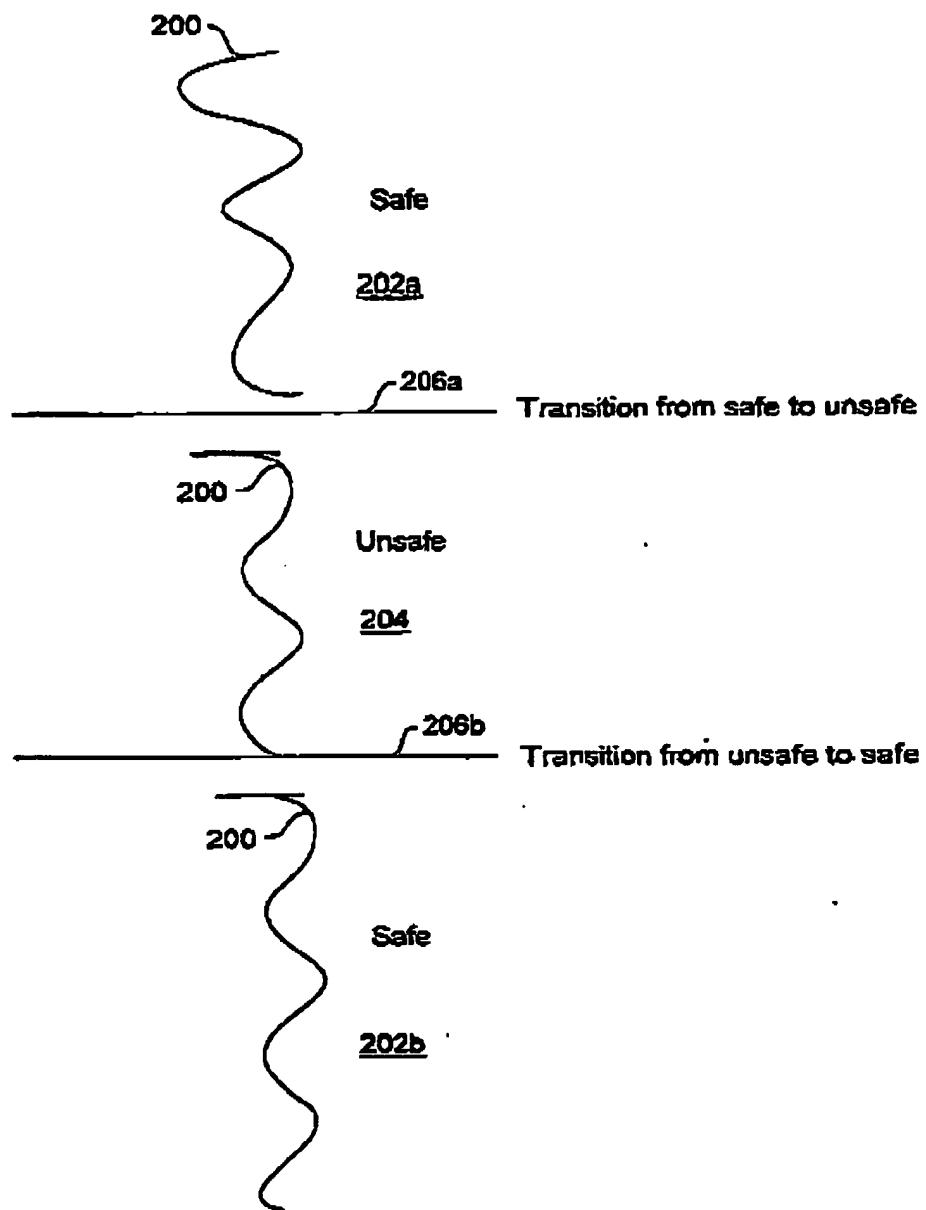


FIG. 2

EP 1 014 264 A2

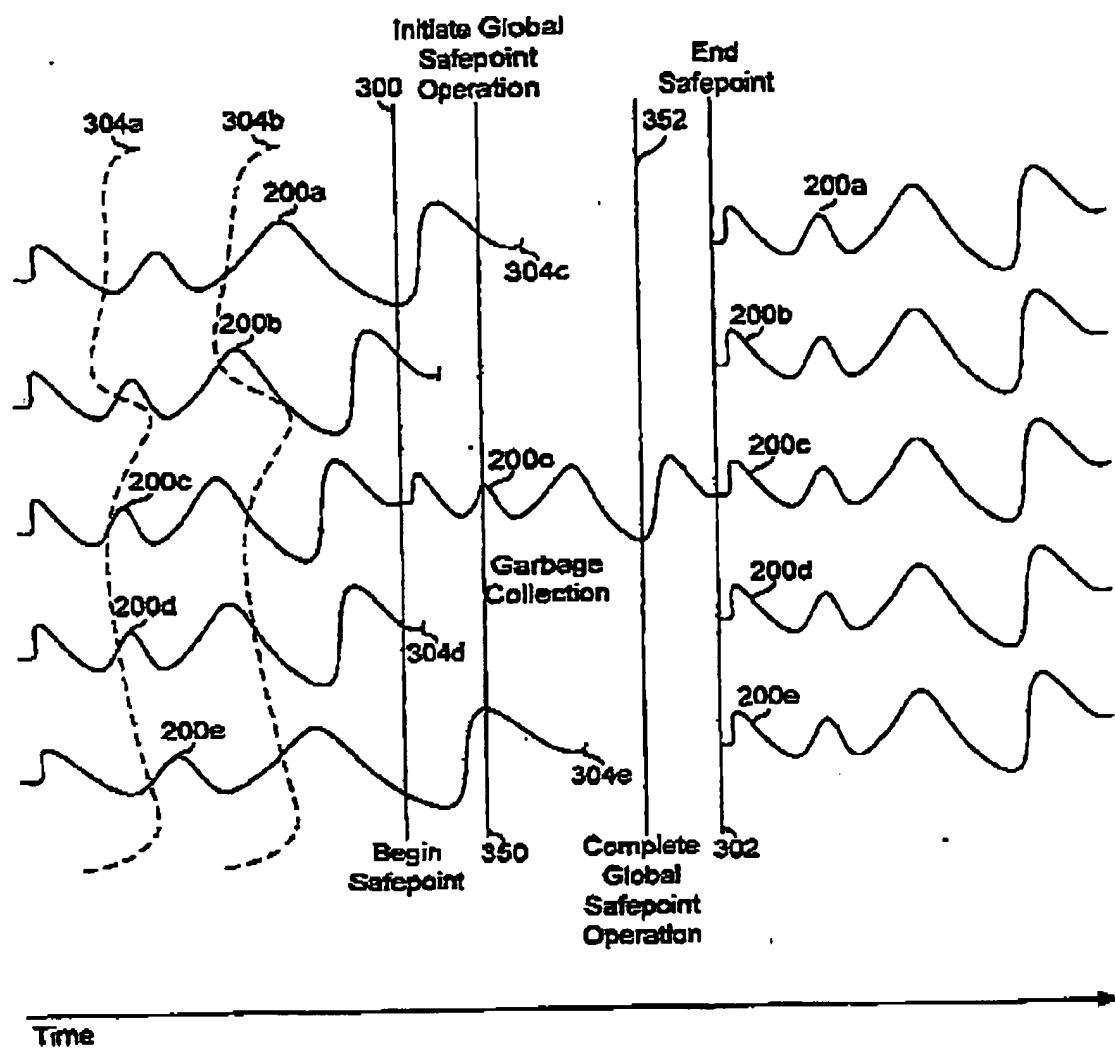


FIG. 3

EP 1 014 264 A2

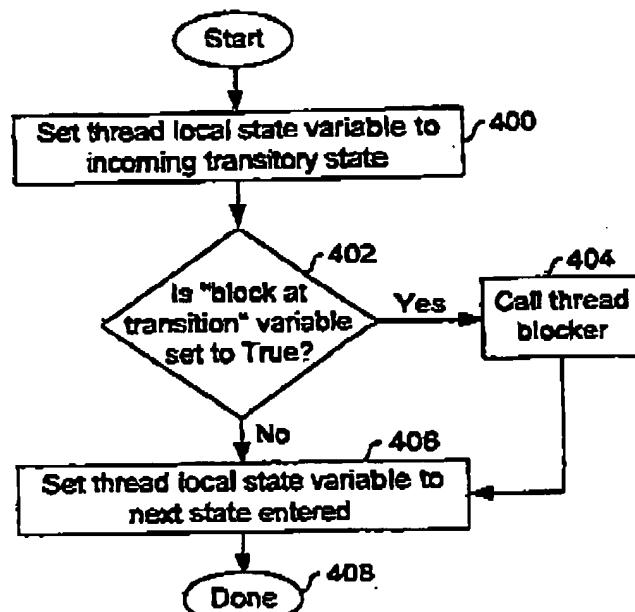


FIG. 4

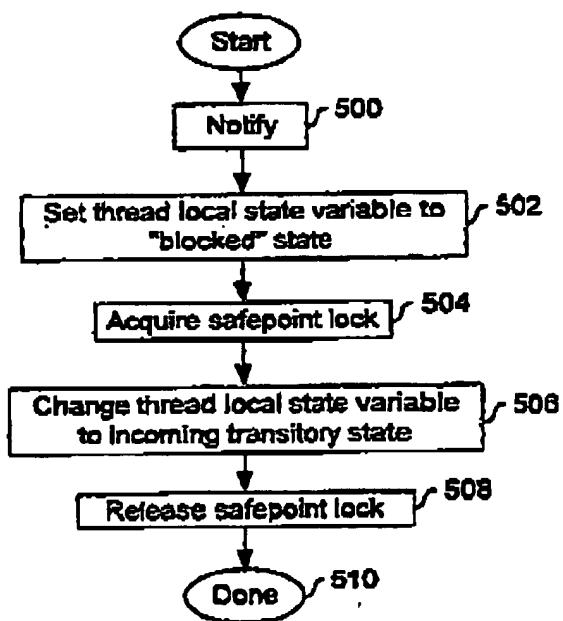


FIG. 5

EP 1 014 264 A2

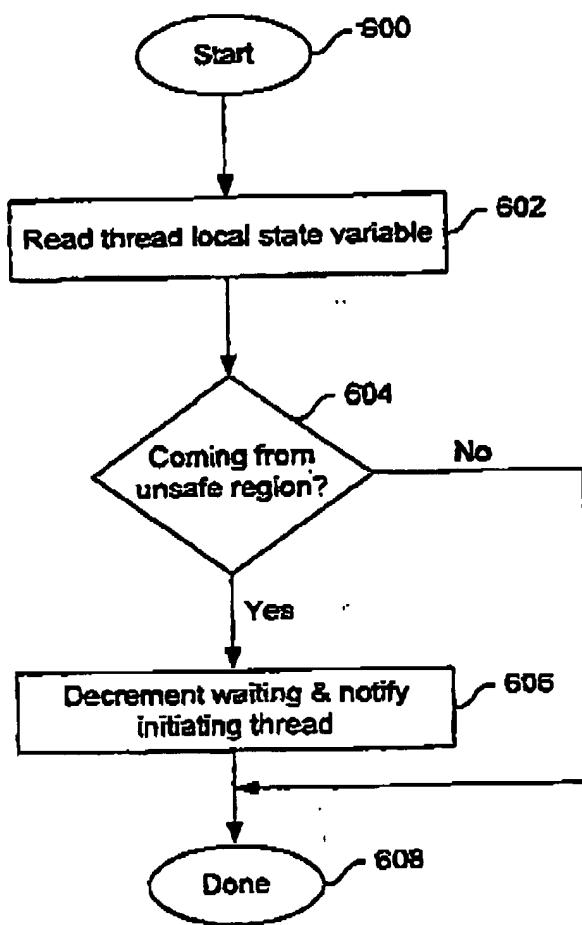


FIG. 6

EP 1 014 264 A2

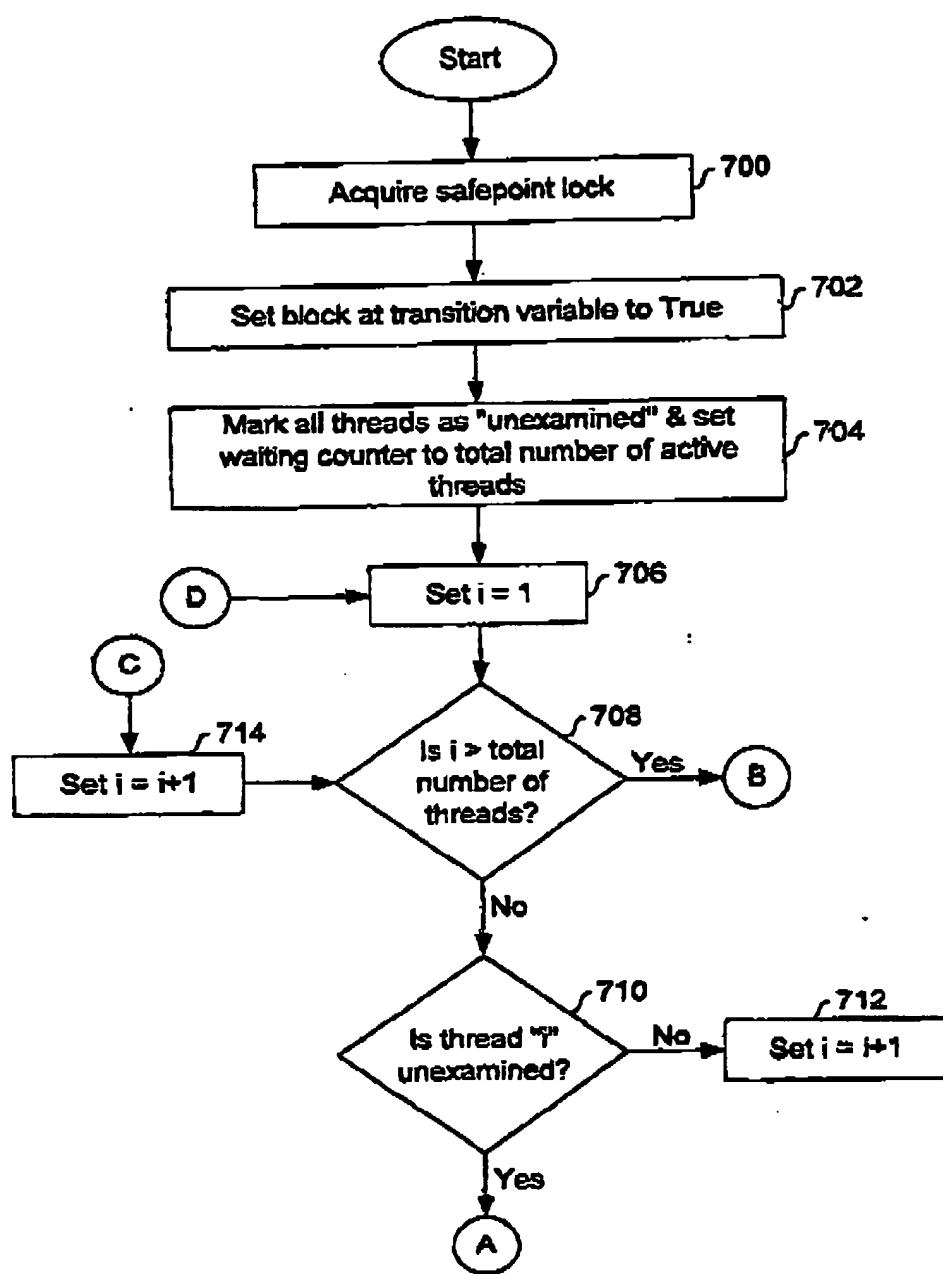


FIG. 7A

EP 1 014 264 A2

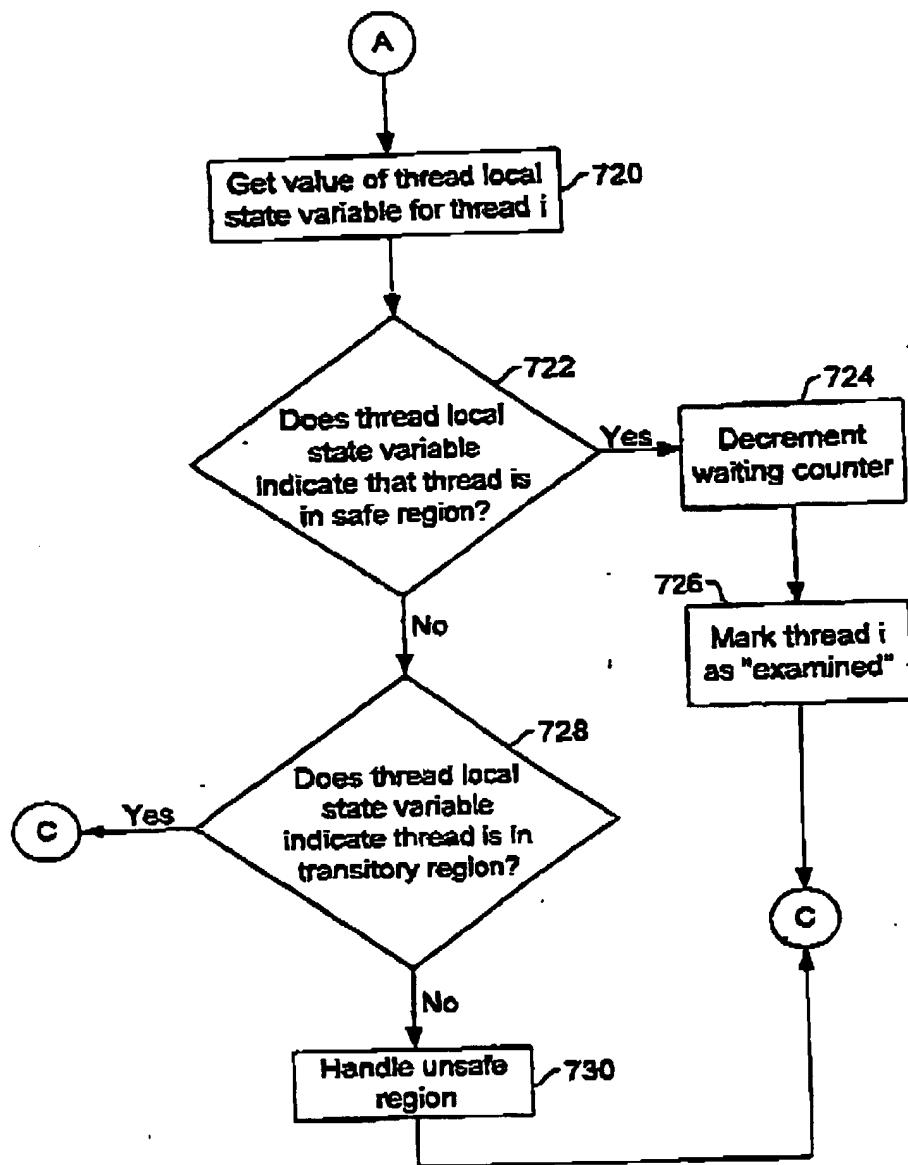


FIG. 7B

EP 1 014 264 A2

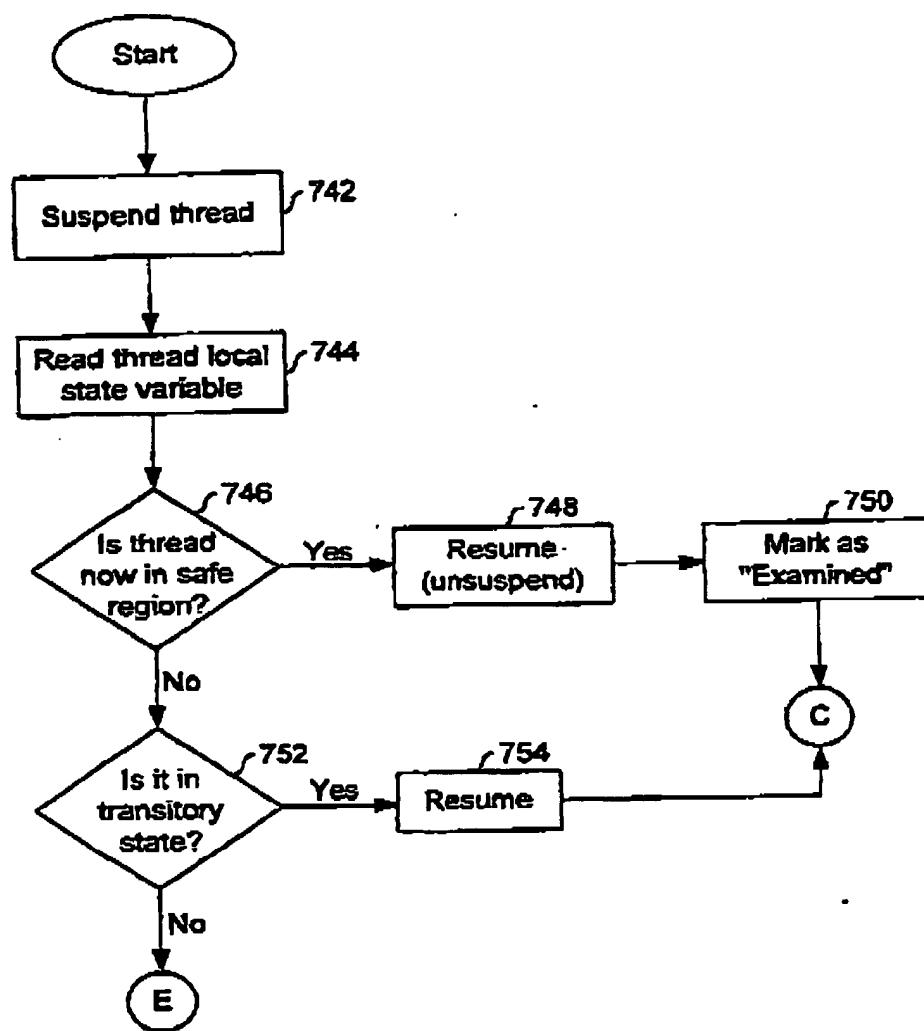


FIG. 7C

EP 1 014 264 A2

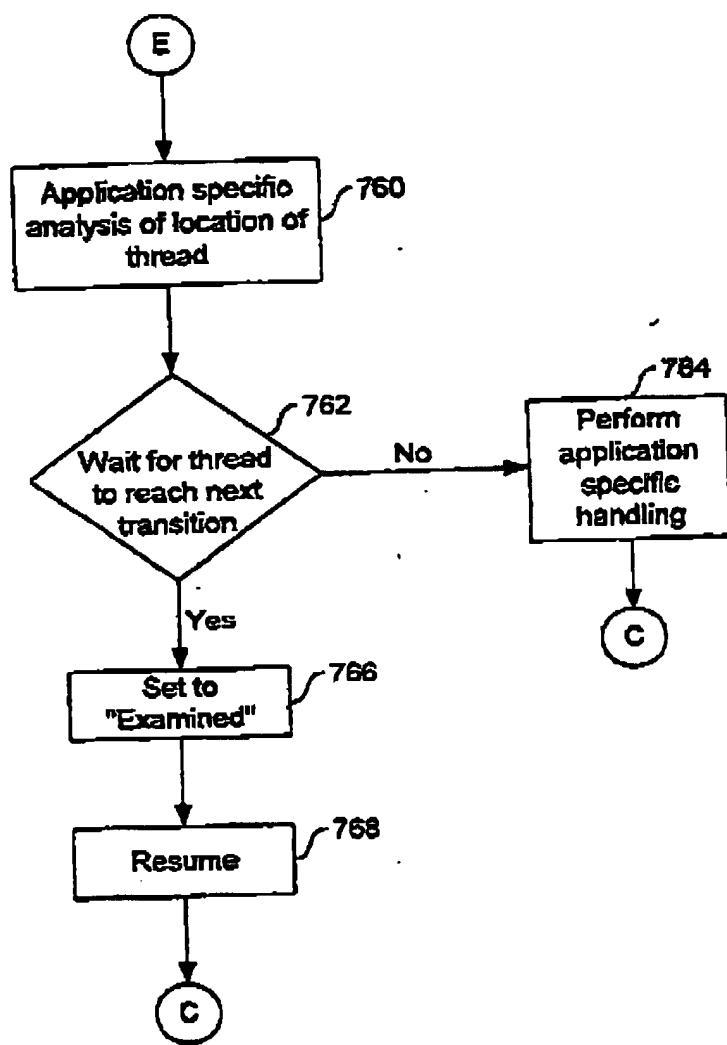


FIG. 7D

EP 1 014 264 A2

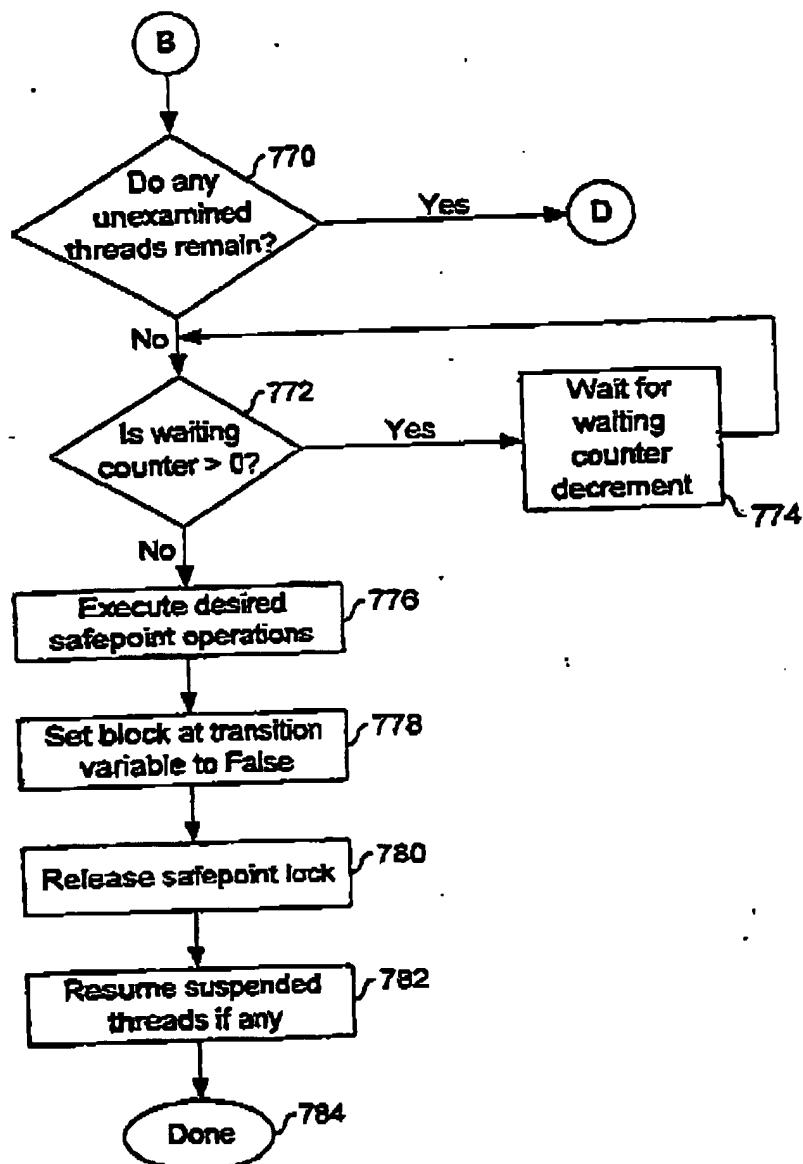


FIG. 7E